Omega 114 (2023) 102745

Contents lists available at ScienceDirect

Omega

journal homepage: www.elsevier.com/locate/omega

Controlling inventories in omni/multi-channel distribution systems with variable customer order-sizes☆

Peter Berling[a,b,c,*], Lina Johansson[b], Johan Marklund[b]

[a] Department of Industrial Economics and Technology Management, NTNU, Postboks 8900, Trondheim, 7491, Norway
[b] Department of Industrial Management and Logistics, Lund University, PO Box 118, Lund, SE-221 00, Sweden
[c] Department of Accounting and Logistics, Linnaeus University, Växjö, SE-351 95, Sweden

ARTICLE INFO

Article history:
Received 24 August 2021
Accepted 15 August 2022
Available online 20 August 2022

Edited by Prof. Benjamin Lev

Keywords:
Inventory
Multi-channel
Multi-echelon
Stochastic
Service level

ABSTRACT

The fast growth of e-commerce and omni/multi-channel retailing brings new challenges for efficient inventory management. One such challenge concerns service differentiation across channels when upstream central warehouses satisfy both direct customer demand and replenishment orders from downstream retailers. Motivated by industry collaboration, we address this issue by developing a combined stock method for control of one-warehouse-multiple-retailer inventory systems with direct customer demand at the central warehouse. The combined stock method, used for service differentiation at the central warehouse, may be described as a critical level policy. The computationally efficient heuristics we present are designed to deal with real-life one warehouse multiple retailer inventory systems characterized by highly variable customer order-sizes, (R, Q) policies at all stock points, and fill rate constraints. A numerical study, including real data from two different companies, illustrates that the heuristics perform well; offering near optimal solutions close to target fill rates, with significant opportunities to reduce total inventory costs compared to existing methods.

1. Introduction

The motivation for this work stems from collaboration with a global inventory management software provider, and the fast growth of e-commerce and omni/multi-channel retailing. An issue of increasing importance for many of the software provider's clients is how to achieve efficient inventory control in large multi-echelon inventory distribution systems with customer demands both at downstream retailer locations and at the upstream central warehouse. The company is therefore interested in finding efficient methods for inventory control of such systems that is conceptually and computationally simple enough to be implemented in practice. More specifically, the methods should determine near optimal reorder points in systems with continuous review (R, Q)-policies, complete back-ordering, fill rate constraints and customer demand with highly variable customer order-sizes.

In this paper we address these issues by considering a one-warehouse-multiple-retailer (OWMR) system where the central warehouse faces demand from two different channels, replenishment orders from the retailers and direct customer demand, typically with very different service requirements. The end customers generally expect high service levels, whereas the multi-echelon inventory literature shows that service to the retailers typically should be relatively low in order to minimize the inventory costs for the entire system, see for example [1,2] and [3]. Service differentiation across these two channels at the central warehouse is therefore important in finding efficient solutions. However, it also complicates the analysis which may be one reason for the limited number of papers on inventory control of these types of multi-echelon systems with direct upstream demand.

One method to differentiate the service level observed both in the literature and in practice is to have a separate inventory (or separate stock) reserved for handling the direct upstream demand at the central warehouse. This can be modelled within a traditional OWMR framework by introducing a separate artificial retailer that serves the upstream demand, see, for example, [4]. This artificial retailer is assumed to replenish its stock from the central warehouse (CW) using an order-up-to S policy and since they are co-located the transportation time from the CW is negligible. We will denote heuristics based on this approach as *separate stock* heuristics. The approach is appealing in many ways as one can apply

☆ Area: Supply Chain Management. This manuscript was processed by Associate Editor Archetti.

* Corresponding author.
E-mail address: lars.p.berling@ntnu.no (P. Berling).

https://doi.org/10.1016/j.omega.2022.102745

established methods for OWMR-systems. However, the performance of these separate stock heuristics deteriorates in systems with fill rate constraints and large variations in customer order-sizes. The problem is that the need for safety stock and thereby order-up-to level at the artificial retailer tend to be overestimated when it is treated as any other retailer. The reason is that the fill rate calculation is based only on the inventory level at the artificial retailer. Our present work contributes to the literature by developing a new *combined stock* heuristic that considers the total amount of inventory currently available at the central warehouse, when determining the critical reservation level S. The *combined stock* includes both the stock reserved at the artificial retailer for the direct upstream demand and the un-reserved general stock at the central warehouse.

This new *combined stock* heuristic can, in principle, be used in conjunction with any method for determining the reorder points in traditional OWMR systems. However, exact solutions procedures, for example, the method in [5], are for computational reasons not a viable option for the real-life problems faced by the software provider's clients. To deal with these situations, an approximate method is required. Our present work focuses on integration of the combined stock heuristic with modified versions of the approximation methods for inventory control of OWMR systems in [6] and [7]. The motivation for building on these approximation methods is that they: (i) are general enough to be directly applied to the considered systems with continuous review, (R, Q) policies and customer demand with highly variable order-sizes, (ii) have shown to perform well, and (iii) are conceptually and computationally simple enough to be implemented in practice. These methods are also partly implemented in the software provider's system. The method in [6] assumes compound Poisson distributed demand, while [7] considers an adjusted normal demand approximation. Both methods use an induced backorder cost heuristic for decomposing the multi-echelon inventory system into $N+1$ single-echelon problems that is solved in a single pass. An added challenge is that the combined stock heuristic requires a modified approach for determining the induced backorder cost associated with the direct customer demand. We suggest and evaluate two such methods; one naïve and one iterative. The latter shows best performance and is used in the proposed BJM heuristic. The former is used for benchmarking purposes in the separate stock heuristic, BM-S, and in the alternative combined stock heuristic, BM-C.

A numerical simulation study, including both real data from two different companies and researcher generated examples, shows that the new BJM-heuristic performs very well. The fill rate targets of the direct customer demand are achieved with high accuracy. For the two test series based on real data, the average difference between the achieved fill rates and the target fill rates are 1.8 percentage points (pp.) and 1.6 pp. respectively. For the researcher generated data, the average deviation is 0.07 pp. when using the BJM-heuristic based on compound distributed demand and −0.68 pp. when using BJM-heuristic based on the adjusted normal demand approximation. At the same time, the average inventory costs in the system are significantly reduced compared to the separate stock benchmark, BM-S. For the test series based on real company data, the average reductions are 5.8 % and 9.7 %, with maximums of 14.3 % and 27.7 %, respectively. For the researcher generated problems, the reduction is on average 9.9 % with a maximum of 16.3 %. Optimization through simulation search also shows that the BJM solutions are near optimal.

The remainder of this paper is organized as follows. Section 2 provides a brief literature review, followed by a model formulation in Section 3. Section 4 derives the proposed heuristics and Section 5 presents the numerical study. Finally, Section 6 summarizes and concludes.

2. Literature review

Omni/multi-channel on-line/off-line retailing is a growing field in practice and in the research literature, e.g., [8–11] and [12]. Much of the literature, though, has focused on a sales and marketing perspective, particularly the literature on omni-channel networks, see [13]. Operational aspects, such as inventory control, which is the focus of our present work, have so far not received the same amount of interest. Exceptions include [14], which investigates the ordering decision for an off-line channel in an omni-channel network under a newsvendor setting. The authors assume that in case of a stock-out, the off-line customers will turn to the on-line channel. The on-line channel is assumed to have an abundance of stock, but there is an increased cost of serving customers through this channel. The reason is the high probability of customer returns when customers cannot see, feel, or try on the product before purchasing it. In [15], the authors extend the model from their previous paper by allowing for Buy-Online-and-Pick-up-in-Store (BOPS), which is a solution where the off-line channel shares its inventory status with the public. This is an approach that empirically has been shown to increase the profit through increased sales, see [16]. However, [9] shows that BOPS may either benefit or hurt the retailer depending on how large the store visiting costs and the on-line waiting costs are.

In [17] and [18] the inventory decisions in a periodic setting for a system with one on-line and one off-line store are investigated. Excess stock at the latter can with additional costs be used to satisfy demand at the former, at the end of each period. This means that these papers are related to the literature on lateral transshipments, see [19] for a review, and to allocation of orders to fulfillment centers for on-line sales, e.g., [20] and [21]. Inventory transshipments are also considered in [22] focusing on assessing product availability in an omni-channel retail network with product substitution. In [17] a single period newsvendor problem is considered, whereas [18] considers a series of such problems. The related literature on inventory control in dual- and multi-channel systems also typically considers either single-period problems, or a series thereof, see, for example, [23,24] and [25]. For a review of this literature, see [26]. Two exceptions are [27] and [28] which consider a continuous review inventory problem with an on-line and an off-line channel under the assumption of exponential lead times and Poisson demand. The stylized lead time assumption allows the authors to solve the problem exactly by total enumeration, a cumbersome method for most real-life problems. All references above, except [23], assume only two out-lets, which seems to be the most common assumption. One can easily argue that the "two-out-lets-single-period" mold, typically used in the literature so far, is not sufficient for the companies studied in our present work, or for any omni/multi-channel network with many retailers experiencing fixed shipping/ordering costs and selling products over multiple periods. The reason is that for all of these inventory systems, it is typically beneficial to order and/or ship in batches and allow the inventory locations to replenish when they are running low on stock rather than with the same fixed intervals.

There is obviously a close connection between our present work and the literature on analysis and inventory control of traditional OWMR systems with demand only at the retailers. Efficient control of such systems under stochastic demand is a challenging problem that has received a lot of research attention over the years. For overviews, see, for example, [1,3,29], and [30]. The latter describes and reviews the literature based on the so called guaranteed service approach for tactical positioning of safety stocks in multi-stage supply chains. This approach is based on the assumption of a bounded demand or that external measures, e.g., overtime, can be used to handle excess demand to always meet the guaranteed service time in a cost efficient manner. This was deemed to not be the case for the software provider's clients, and the real cases considered in our present work.

The literature on inventory control of OWMR systems contains both exact solution procedures and a wide range of approximation methods subject to various assumptions and constraints. However, relatively few of these multi-echelon methods have found their way into practical applications, the guaranteed service approach is a prominent exception. At least part of the explanation is that many are based on quite restrictive assumptions (e.g., identical retailers, specific demand distributions, no order quantities etc.) or that they are computationally demanding.

Our work is also related to the literature on critical level policies in single-echelon systems with different demand classes. To the best of our knowledge, the critical level policy was first suggested and analyzed by Veinott [31] and early work in the field includes [32,33], and [34]. More recent work in that area include, for example, [35–37] and [38]. An overview of the literature on these and other types of rationing policies is found in [39]. An approach, conceptually similar to ours is [40] where artificial retailers are introduced to differentiate the service and transform the single-echelon problem into a divergent multi-echelon problem with FCFS clearing of backorders. Compared to our present work important differences are that [40] considers a single-echelon system with base-stock ordering, whereas we consider a multi-channel network with batch ordering at the retailers and the central warehouse. Furthermore, we focus on developing an efficient and easy-to-implement heuristic for real demand with variable customer order-sizes, whereas [40] presents a method for an exact evaluation of the policy limited to Poisson demand.

It is apparent that a FCFS clearing policy is not optimal in a single-echelon system with multiple demand classes. Instead, a priority clearing mechanism should be used where backorders for high priority customers are cleared, and the stock reserved for them are refilled, before tending to backorders from customers with lower priority. However, invoking such a policy renders the problem analytically intractable. As an approximation, [38] proposes a threshold clearing policy that shows good results. In a multi-echelon system, the optimal clearing mechanism at the central warehouse is further complicated by the fact that all units will eventually satisfy end customer demands. Some of these end customers might have a higher priority than the ones who place their orders directly at the warehouse. Finding an optimal clearing policy is beyond the scope of our present work and in accordance with [40] and [4], we assume that backorders are cleared in a FCFS sequence.

3. Model formulation

As mentioned above, we consider an omni/multi-channel OWMR inventory system with end customer demand both at the central warehouse and at the N non-identical retailers. The end customer demand at the central warehouse is referred to as direct upstream demand. The model has been developed in close dialogue with the aforementioned software provider to assure it offers a valid representation of the reality their clients face.

A majority of the software provider's clients are aftermarket service providers. They typically stock many different SKU's, both spare parts and consumables, for which the demand is often low and lumpy (i.e., intermittent with large differences in customer order-sizes) with large variance to mean rations for the demand per time unit. Also, for items with higher demand, variable customer order-sizes is a common and challenging characteristic. The retailers (or dealers) are serving end customers in different geographical regions which means that for practical reasons customers generally cannot choose which retailer location to go to. Apart from replenishing the retailers, the central warehouse may also directly serve end customers in the region where it is located to avoid establishing a separate retail location there. The direct upstream demand may also come from customers in regions far from any of the retail locations, or from independent retailers, which do not belong to the centralized OWMR system.

Motivated by these system characteristics, and particularly the highly variable customer order-sizes, our main approach is to model the customer demand (both at the retailers and the direct upstream demand at the warehouse) as independent compound Poisson processes. Thus, for each location i the customer demand process is specified by the customer arrival rate λ_i, and the discrete stochastic order-size, O_i. We also consider an alternative approximation model with adjusted normal demand, defined in accordance with [7]. This latter model is computationally attractive for high demand items and can approximate demand that is not compound Poisson (e.g., with a variance to mean ratio less than one).

As mentioned above, for modeling purposes, and in accordance with [4], we assume that the direct upstream demand at the central warehouse is handled by an artificial retailer, denoted by index $N+1$, corresponding to a reservation stock at the central warehouse. As a result, the total stock at the central warehouse can be split in two parts; the reservation stock dedicated to serving the direct upstream demand, and the general warehouse stock, denoted by index 0, that is used to replenish all retailers including the artificial retailer.

The lead time from the outside supplier to the central warehouse, L_0, and the transportation times from the central warehouse to the regular retailers, l_i for $i = 1, \dots, N$ are assumed to be positive and constant. These are standard assumptions in the OWMR inventory literature (see, for example [1,3] and [41]). It is also consistent with what is currently assumed in the software provider's system and thereby with the real data we have evaluated. However, it is undoubtedly so that in many real systems the lead time from the outside supplier and/or the transportation times from the warehouse to the retailers may be subject to significant uncertainties. In these situations, our model can still be used after applying standard approximation techniques for dealing with stochastic lead times in inventory systems, e.g., replacing the stochastic lead with its mean as in the famous METRIC model [42,43], or to use the mean and variance of the lead time to determine the mean and variance of the lead time demand, and then fit an appropriate distribution to this first two moments (see, for example, [7] and [41]). It should be noted that the transportation time from the central warehouse to the artificial retailer (which is integrated with the warehouse), l_{N+1}, is zero. Moreover, the lead times (L_i, $i = 1, \dots, N, N+1$) for retailer replenishment orders are stochastic because there may be shortages of general warehouse stock.

The central warehouse and the N retailers are assumed to apply continuous review installation stock (R, nQ) policies to replenish their inventory. Thus, an order of size Q is placed when the inventory position (stock on hand + outstanding orders - backorders) falls to or below the reorder point R. This policy assumption is motivated by its wide-spread use among the software provider's clients. The inventory management system is transaction based and can react immediately when transactions are recorded in accordance with a continuous review policy. From an analytical point of view, the continuous review assumption simplifies the policy evaluations and may adequately approximate periodic review systems with short review periods, for instance by increasing the mean replenishment lead-time by half of the review period, see [41]. It should be noted that many of the software provider's clients still prefer to place orders in the system at given time instances, typically once a day. For such short review periods the continuous review assumption (which is currently used in the software provider's system when determining the reorder points) has proven to work well as an approximation. Similar to most of the literature on stochastic OWMR inventory systems, see [1,3] and references therein, our objective is to optimize the reorder points for given order quantities. The latter may be determined using a deterministic model, see, for example, [44]. This type of approach is motivated by the observation that using deterministic lot sizing methods in a stochastic environment tend to have a small impact on the expected cost as long as the reorder points are adjusted appropriately, see for example [45] and [46]. Moreover, in practice the choice of order quantities are often constrained by load carrier and package sizes, rendering a limited number of alternative order quantities that can easily be evaluated by repeated use of our proposed heuristics.

The artificial retailer is integrated with the central warehouse and replenishes from the general warehouse stock according to a continuous review $(S-1, S)$ policy (or equivalently an $(R_{N+1}, 1)$ policy), where the order-up-to level S can be interpreted as a critical reservation level for the combined stock at the central warehouse. All unsatisfied demand is backordered and satisfied in a FCFS manner. This means there is no difference in priority between the retailer orders satisfied from the general warehouse stock (including those placed by the artificial retailer). As mentioned above, the majority of the software provider's clients are aftermarket service providers that sell spare-parts and consumables, which are often difficult to substitute (due to warranty issues and limited availability of white label alternatives for equipment specific parts). This means that their customers typically are willing to wait for the products to become available, in accordance with the full backordering assumption.

For companies selling standard products the probability of lost sales or spill over to other products when there is a shortage might be significant. However, one can argue that if the fill rate is high a complete backorder model is still a good approximation as the probability of a shortage is low, see [47].

The considered costs are holding costs per unit and time unit at the central warehouse for the general stock, h_0, and the reservation stock, h_{N+1}, and holding costs per unit and time unit at all regular retailers, h_i for $i = 1, \dots, N$. In practice, typically $h_{N+1} = h_0$ so this is what we assume in the analysis, although this assumption is by no means necessary. All retailers (including the artificial retailer) operate under fill rate constraints with specified fill rate targets, γ_i^* for $i = 1, \dots, N, N+1$. The fill rate is defined as the fraction of demand that can be satisfied directly from stock on hand. We focus on (demand) fill rate constraints because this is what the software provider and its two clients that we have data from use. However, it is straight forward to extend the analysis and the proposed heuristics to, for example, order-line fill rate constraints (i.e., the fraction of complete customer orders of a single item satisfied directly from stock), or backorder costs per unit and time unit.

IP_i	–	inventory position at location i $(i = 0, 1, \dots, N, N+1)$
IL_0	–	inventory level of the general stock at the central warehouse
IL_{N+1}	–	inventory level of the reservation stock at the central warehouse (corresponding to the inventory level of the artificial retailer)
IL_{cw}	–	combined inventory level at the central warehouse, $IL_{cw} = IL_0 + IL_{N+1}$
IL_i	–	inventory level at retailer i $(i = 1, \dots, N)$
γ_{N+1}	–	fill rate for the direct upstream demand at the central warehouse
γ_i	–	fill rate at retailer i $(i = 1, \dots, N)$
γ_{N+1}^*	–	target fill rate for the direct upstream demand at the central warehouse
γ_i^*	–	target fill rate for the demand at retailer i $(i = 1, \dots, N)$
μ_i	–	expected demand per time unit at retailer i $(i = 1, \dots, N)$
μ_{N+1}	–	expected direct upstream demand (i.e., at the artificial retailer) per time unit
μ_0	–	expected total demand at the central warehouse per time unit, $\mu_0 = \sum_{i=1}^{N+1} \mu_i$
O_i	–	stochastic customer order-size at retailer i $(i = 1, \dots, N, N+1)$
$f_{O_i}(d)$	–	probability that $O_i = d$ $(d = 1, \cdots, d_{max})$
R_0	–	reorder point for the general stock at the central warehouse
S	–	critical reservation level of the combined stock at the central warehouse (corresponding to the base-stock level at the artificial retailer)
R_{cw}	–	reorder point for the combined stock at the central warehouse, $R_{cw} = R_0 + S$
Q_0	–	order quantity for replenishments placed by the central warehouse with the outside supplier
R_i	–	reorder point at retailer i $(i = 1, \dots, N)$
$\boldsymbol{R}$	–	$(R_1, R_2, \dots, R_N)$
Q_i	–	order quantity at retailer i $(i = 1, \dots, N)$
z^+	–	$\max(z, 0)$
z^-	–	$\max(-z, 0)$

The objective is to minimize the expected total holding cost per time unit, TC, by optimizing R_0, S and R_i for $i = 1, \dots, N$, subject to fill rate constraints at the retailers, and for the direct upstream demand at the central warehouse

$$\begin{aligned} \min\ & TC(R_0, S, \boldsymbol{R}) = h_0 E[IL_{cw}^+] + \sum_{i=1}^{N} h_i E[IL_i^+] \\ \text{s.t. } & \gamma_{N+1}(R_0, S) \geq \gamma_{N+1}^* \\ & \gamma_i(R_0, R_i) \geq \gamma_i^* \quad \forall\ i = 1, \dots, N. \end{aligned} \tag{1}$$

This is a complex problem to solve as all the fill rates depend on the reorder point for the general warehouse stock, R_0, as well as the local reorder point, R_i, $i = 1, \dots, N$ and S, in a non-trivial manner.

4. Proposed heuristics

As previously stated, the goal of our present work is to find an efficient heuristic for optimizing the critical reservation level S for the direct upstream demand at the central warehouse and the reorder points R_i, for $i = 0, 1, \dots, N$, in the considered system. Our approach can be divided in two main steps. The first step is to obtain the combined stock heuristic for optimizing the critical reservation level S. The objective is to find the smallest S that satisfies the fill rate requirements for the direct upstream demand at the warehouse, given a reorder point, R_0, for the general warehouse stock. The second step is to integrate the combined stock heuristic in an efficient method for inventory control of the entire OWMR system by optimizing the reorder points R_0 and R_i, for $i = 1, \dots, N$. Here we have chosen to modify the approximation methods in [6] and [7] because of the good performance they exhibit in earlier studies, and their flexibility and ability to deal with the reality faced by the software provider's clients, including the two companies we have real data from. However, it should be noted that the combined stock heuristic may be used in conjunction with any method for regular OWMR systems to control multi-channel distribution systems with inventory reserved for direct upstream demand. In coherence with the real demand data our main approach is a heuristic based on a modification of the method in [6] that assumes compound Poisson demand. We also propose an alternative heuristic by modifying the method in [7] which is based on a normal demand approximation. This method is computationally simpler and represents a more general approximation method that can deal also with demand that is not compound Poisson.

In the following, Section 4.1 presents the combined stock heuristic for compound Poisson demand, and Section 4.2 shows how this heuristic is modified to adjusted normal demand. Section 4.3 then explains how the approximation models in [6] and [7], jointly referred to as BM, are adapted and integrated with the combined stock heuristics to arrive at the two versions of a complete heuristic BM-C and BJM. This section also defines the separate stock heuristic BM-S that is used for benchmarking purposes.

4.1. The combined stock heuristic for compound Poisson distributed upstream demand

The purpose of the combined stock heuristic is to determine the smallest critical reservation level S for the combined warehouse stock that satisfies the fill rate constraint for the direct upstream demand, given the reorder point R_0 for the general warehouse stock. Our approach is based on approximating the probability mass function (pmf) for the combined inventory level of the reservation stock and the general stock at the central warehouse, given S and R_0. From a modeling perspective, we assume that an artificial retailer with base-stock level S and zero transportation time from the warehouse handles the direct customer demand at the warehouse. This means that as soon as a demand occurs, the artificial retailer replenishes from the general warehouse stock, which also satisfies replenishment orders from the N regular retailers in a FCFS manner. The combined stock policy prescribes that when a customer arrives and demands d units, it will be satisfied by the combined stock on hand at the central warehouse IL^+_{CW}, which is the sum of the stock on hand at the artificial retailer ($\leq S$) and the available general warehouse stock. Thus, even if $d > S$ the demand may be fully satisfied if $d - S$ units are available in the general warehouse stock. Recall that d is a realization of the stochastic demand size O_{N+1}. This is in contrast to the separate stock policy (see [4]), which implies that the fill rate is calculated based solely on the inventory level at the artificial retailer. Thus, in the separate stock approach, at most S demanded units are considered to be satisfied immediately from stock on hand even if there is general stock available at the central warehouse.

The added challenge in using the combined stock policy is to determine the probability distribution of the combined stock on hand, IL^+_{CW}. The exact distribution is unknown but inherently complex, and for computational reasons most likely infeasible to use in the real systems we consider. Thus, we focus on obtaining an efficient approximation.

First, note that the inventory level of the combined stock, IL_{CW}, is equal to $IL_0 + S$ when the inventory level of the general warehouse stock is positive, i.e., $IL_0 > 0$. When $IL_0 \leq 0$, the combined stock on hand IL^+_{CW} is equal to IL^+_{N+1} and all available inventory is reserved for the direct upstream demand. The probability distribution of the combined stock on hand $IL_{\mathrm{CW}} \geq 0$ (for given R_0 and S) may then be determined as

$$P(IL_{\mathrm{CW}} = j|R_0, S) = \begin{cases} P(IL_0 = j - S) & \text{if } j > S \\ P(IL_0 \leq 0)P(IL_{N+1} = j|IL_0 \leq 0) & \text{if } 0 \leq j \leq S. \end{cases} \tag{2}$$

Let $D_0(t)$ be the total demand at the central warehouse during t time units. Note that $D_0(t)$ incorporates the impact of retailer order quantities and customer order-sizes in the compound Poisson demand processes at the retailers. The probability distribution of $D_0(t)$ can be determined, for example, using the methods in [6]. The probability distribution for IL_0 and $P(IL_0 \leq 0)$ (given R_0 and Q_0) can then be obtained as

$$P(IL_0 = j) = \frac{1}{Q_0} \sum_{k=\max(R_0+1,j)}^{R_0+Q_0} P(D_0(L_0) = k - j), \quad 1 \leq j \leq R_0 + Q_0 \tag{3}$$

and

$$P(IL_0 \leq 0) = 1 - \sum_{j=1}^{R_0+Q_0} P(IL_0 = j), \tag{4}$$

The probability distribution for IL_{N+1} is more difficult to determine, as the replenishment lead time L_{N+1} is stochastic because the general warehouse stock may be depleted. To arrive at an efficient approximation, first note that a delay occurs only when there is no general stock on hand, i.e., $IL_0 \leq 0$. Let $\hat{L}$ denote the expected lead time for units ordered by the artificial retailer that experience a delivery delay (and thus a lead time greater than zero), because there is no general warehouse stock available. Focusing only on these situations, i.e., when $IL_0 \leq 0$, and assuming (as an approximation similar to the METRIC approach [42,43]) that the lead time is constant and equal to its mean, we have $IL_{N+1} = S - D_{N+1}(\hat{L})$, where $D_{N+1}(t)$ denotes the stochastic demand at the artificial retailer during t time units derived from the compound Poisson Process. The pmf for the stock on hand at the artificial retailer is then

$$P(IL_{N+1} = j|IL_0 \leq 0) = P(D_{N+1}(\hat{L}) = S - j). \tag{5}$$

To estimate $\hat{L}$, we assume that the lead time L_{N+1} follows a two point distribution such that it is zero when there are items available in the general warehouse stock, and $\hat{L}$ when there is a delay. The probability that there is a delay is denoted by α. Furthermore, we assume that the mean lead time $E[L_{N+1}]$ is equal to the expected delay per unit delivered from the general warehouse stock, $\bar{L}$, determined using Little's law, $\bar{L} = E[IL_0^-]/\mu_0$. Given these assumptions, we have $\bar{L} = \alpha\hat{L} + (1-\alpha)\cdot 0$ and thereby

$$\hat{L} = \frac{\bar{L}}{\alpha}. \tag{6}$$

Note that for the situations we consider, $\bar{L}$ is not necessarily the correct average delay for the units delivered to a certain retailer. However, it tends to be a robust approximation with good performance, particularly if partial deliveries are applied, see for example [2].

The probability α may be estimated in different ways. We propose to set $1-\alpha$ equal to the ready rate for the general stock at the central warehouse, RR_0, which by definition is the probability that $IL_0 > 0$. Note that this corresponds to the proportion of time that at least some part of an upstream demand order is satisfied without a delay. This give us the following estimation of $\hat{L}$,

$$\hat{L} = \frac{\bar{L}}{(1 - RR_0)}. \tag{7}$$

Another possibility to estimate the probability $1-\alpha$ is to use the fill rate at the general warehouse stock for demands from the artificial retailer. This alternative has been evaluated numerically, but the performance of this estimation is inferior when it comes to fulfillment of the fill rate targets of the upstream demand. The ready rate is also simpler to calculate for more variable demand patterns, compared to the fill rate. Because the ready rate is greater than or equal to the fill rate, using the former gives us a larger estimate for $\hat{L}$, suggesting that a larger lead time variability is incorporated into the estimate of the combined stock on hand at the central warehouse. Note that, if the demand at the central warehouse is Poisson distributed, the fill rate and ready rate coincide and $\bar{L}$ is the correct expected lead time, see for example [41] pp. 79, 203. The fill rate for the direct upstream demand is easily obtained from (8) once the distribution of the combined stock on hand $P(IL_{\mathrm{CW}} = j|R_0, S)$ $\forall j \geq 0$ is determined.

$$\gamma_{N+1}(R_0, S) = \frac{\sum_{d=1}^{d_{\max}} \sum_{j=1}^{R_0+S+Q_0} \min(j, d) f_{O_{N+1}}(d) P(IL_{\mathrm{CW}} = j|R_0, S)}{\sum_{d=1}^{d_{\max}} d\, f_{O_{N+1}}(d)} \tag{8}$$

The objective function in (1) stipulates that the reservation level S should be set as low as possible while still fulfilling the target fill rate, γ^*_{N+1}, for the direct demand. It is straightforward to determine the smallest S that satisfies the target fill rate by increasing S from zero until $\gamma_{N+1}(R_0, S) \geq \gamma^*_{N+1}$.

4.2. The combined stock heuristic for adjusted normal demand

The combined stock heuristic for adjusted normal demand offers a computationally simpler alternative to determine S. It may also be used in situations where demand is not compound Poisson (e.g., the variance to mean ratio of the demand per time unit is less than one). The method is based on assuming that the lead time demand is normally distributed and accounts for the variable customer order-sizes by an undershoot adjustment of the reservation level S.

As before, IL_{cw} equals $IL_0 + S$ when $IL_0 > 0$, and it equals IL_{N+1} when $IL_0 \leq 0$. Under the assumption of normally distributed lead time demand, the fill rate for the direct upstream demand is equal to the ready rate, RR_{cw}, i.e., the probability of positive stock on hand, see for example [41] pp. 79, 82. This means

$$\begin{aligned}\gamma_{N+1}(R_0, S) &= RR_{cw} = P(IL_{cw} > 0)\\ &= P(IL_0 > 0) + P(IL_0 \leq 0) \cdot P(IL_{N+1} > 0 | IL_0 \leq 0)\\ &= 1 - P(IL_0 \leq 0) \cdot P(IL_{N+1} \leq 0 | IL_0 \leq 0). \end{aligned} \quad (9)$$

Consequently, it remains to determine the probability $P(IL_0 \leq 0) \cdot P(IL_{N+1} \leq 0 | IL_0 \leq 0)$, γ_{N+1} then follows from (9). Using the same lead-time approximation as before, i.e., that the lead time is constant and equal to $\hat{L}$ when $IL_0 \leq 0$, it is straight forward to calculate $P(IL_{N+1} \leq 0 | IL_0 \leq 0)$ under the normal demand assumption disregarding the customer order-sizes using (10).

$$\begin{aligned}P(IL_{N+1} \leq 0 | IL_0 \leq 0) &= 1 - RR_{N+1}(\tilde{R}_{N+1} | IL_0 \leq 0)\\ &= \sigma'_{N+1}\left(G\left(\frac{\tilde{R}_{N+1} - \mu'_{N+1}}{\sigma'_{N+1}}\right) - G\left(\frac{\tilde{R}_{N+1} + 1 - \mu'_{N+1}}{\sigma'_{N+1}}\right)\right), \end{aligned} \quad (10)$$

where, $\tilde{R}_{N+1}$ is the reorder point, $\mu'_{N+1} = \mu_{N+1}\hat{L}$, $\sigma'_{N+1} = \sigma_{N+1}\sqrt{\hat{L}}$ and G is the general loss function given a normally distributed lead-time demand and that all orders are triggered at the reorder point, (see, e.g., [41] pp. 76-77, 82).

However, in contrast to a truly normally distributed demand, the variable customer order-sizes will lead to a stochastic undershoot of the set reorder point $R_{N+1} = S - 1$. That is, the realized reorder point, i.e., the inventory position when the order is triggered, will for some orders be $\tilde{R}_{N+1} = R_{N+1} - u$ rather than R_{N+1} and this must be accounted for when determining $P(IL_{N+1} \leq 0 | IL_0 \leq 0)$ and the reservation level S. To exemplify, let us consider two situations when a customer arrives and requests 3 units when: (i) $IL_0 = 1$, i.e., $IL_{cw} = S + 1$, and (ii) $IL_0 < 0$, i.e., $IL_{cw} = IL_{N+1}$. In (i), $IL_{cw} = 1 + S - 3 = S - 2$ and $IL_0 = -2$ after the demand has occurred. In this scenario, $IL_0 \leq 0$ and IP_{N+1} will drop to $S - 2$ before a replenishment order is triggered, i.e., there is an undershoot of one unit ($u = 1$) of the set reorder point $S - 1$. In (ii), IP_{N+1} drops from S to $S - 3$ before a replenishment order is triggered, corresponding to an undershoot of 2 units.

To account for the undershoot caused by the variable customer order-sizes when calculating the fill rate under the normal demand approximation we determine the distribution of the realized reorder points, or equivalently the undershoot when $IL_0 \leq 0$ just after the customer demand has been satisfied. To obtain the probability of an undershoot of u units in these scenarios, $P(U = u | IL_0 \leq 0)$, we first sum the probabilities of the customer demand sizes, O_{N+1}, that will cause an undershoot of u units, i.e., $P(O_{N+1} = IL_0 + u + 1)$ for $IL_0 > 0$ before the demand occurs, and $P(O_{N+1} = u + 1)$ otherwise (see examples (i) and (ii) above). This sum is then divided with the probability that a customer demand results in $IL_0 \leq 0$. The resulting expression is

$$P(U = u) = \frac{\sum_{j=0}^{R_0+Q_0} P(IL_0^+ = j) P(O_{N+1} = j + 1 + u)}{\sum_{u=0}^{u_{\max}} \sum_{j=0}^{R_0+Q_0} P(IL_0^+ = j) P(O_{N+1} = j + 1 + u)}, \quad (11)$$

where $u_{\max}$ is equal to the largest undershoot, which is equal to $d_{\max} - 1$.

Combining (10) and (11), we attain the following expression for the fill rate of the direct upstream demand when accounting for the undershoot

$$\begin{aligned}\gamma_{N+1}(R_0, S) &= 1 - P(IL_0 \leq 0) \cdot P(IL_{N+1} \leq 0 | IL_0 \leq 0)\\ &= 1 - P(IL_0 \leq 0) \sum_{u=0}^{u_{\max}} P(U = u)\\ &\quad \times (1 - RR_{N+1}(S - 1 - u | IL_0 \leq 0)), \end{aligned} \quad (12)$$

since the realized reorder point with an undershoot of u units is $S - 1 - u$ and an alternative interpretation of (12) is

$$\gamma_{N+1}(R_0, S) = P(IL_0 > 0) + P(IL_0 \leq 0) \sum_{u=0}^{u_{\max}} P(U = u) FR(S - 1 - u), \quad (13)$$

where $FR(S - 1 - u)$ is the demand fill rate for the reorder point $S - 1 - u$ when demand is $\mathcal{N}\left(\mu_{N+1}\hat{L}, \sigma_{N+1}\sqrt{\hat{L}}\right)$ distributed. Note that under the normal demand approximation the fill rate is 1 for $IL_0 > 0$. Using (12) or (13) S may be adjusted to satisfy the fill rate constraint for the upstream demand, γ^*_{N+1}.

4.3. The OWMR multi-channel heuristics BJM, BM-C and BM-S

As explained above, the OWMR multi-channel heuristics BJM and BM-C are obtained by using the new combined stock heuristics together with modified versions of the BM methods for control of traditional OWMR systems in [6,7]. The BM methods can be described in terms of a five step procedure where the details of the steps are different depending on the demand assumption. In the following, we will augment and modify this procedure to define and explain the details of the BJM and BM-C heuristics as well as the separate stock heuristic BM-S. The latter fits within the first five steps, while BJM and BM-C requires an augmented sixth step. Before going through the procedure, we first note that applying the BM approach to our omni/multi-channel system entails decomposing the multi-echelon problem into $N + 2$ coordinated single-echelon problems that are computationally easy to solve. The decomposition is achieved by introducing an induced backorder cost, β, at the central warehouse, that captures how the retailers are affected by the reorder point R_0 at the central warehouse.

Step 1: Following the BM approach, the induced backorder cost for the warehouse, β, is estimated as a demand weighted average of the induced backorder costs associated with the retailers. In our case this includes both the N regular retailers (with induced backorder costs β_i for $i = 1, \ldots, N$) and the artificial retailer (with induced backorder cost β_{N+1}) according to

$$\beta = \frac{\mu_{N+1}}{\mu_0}\beta_{N+1} + \sum_{i=1}^{N} \frac{\mu_i}{\mu_0}\beta_i. \quad (14)$$

For the regular retailers ($i = 1, \ldots, N$), β_i is determined according to the BM methods. Thus, by using the same tables and closed form expressions, after estimating a backorder cost p_i for each regular retailer that corresponds to its target fill rate, γ_i^* as

$$p_i = \frac{\gamma_i^*}{1 - \gamma_i^*} h_i. \quad (15)$$

Note that (15) is a correct expression for continuous demand, i.e., the smallest reorder point satisfying the fill rate constraint is the one that minimizes the expected cost if and only if the backorder cost per unit and time is p_i. For discrete demand (15) is an approximation because a larger range of backorder costs may render the same reorder point as when a fill rate constraint is used, see, e.g., [41] pp. 86-87 for details.

The induced backorder cost β_{N+1} associated with the artificial retailer requires a different approach as the transportation time, l_{N+1}, is 0 and the available tables and expressions in [2] are only valid for positive transportation times. We therefore suggest and evaluate two new methods, one naïve and one iterative.

The naïve method used in the BM-C and BM-S heuristics, sets β_{N+1} equal to the backorder cost per time unit for the direct demand, p_{N+1}, estimated from (15). The cost per time unit for a delay at the central warehouse cannot exceed the backorder cost per time unit at the retailer, so p_{N+1} may be seen as an upper bound for β_{N+1}. Thus, the naïve approximation $\beta_{N+1} = p_{N+1}$ tends to overestimate the correct value for β_{N+1}. This may lead to a value of R_0 that is too high and too much general warehouse stock.

The iterative method, used in the BJM heuristic, attempts to find a better estimate of β_{N+1} by applying a modified version of iterative procedure in [48]. More precisely, starting with the naïve estimate $\beta_{N+1} = p_{N+1}$, the procedure is adapted to search for a lower estimate of β_{N+1} while β_i for $i = 1, \ldots, N$ are fixed at the values determined according to the BM methods. Further details regarding the iterative search procedure are provided in Appendix B. The numerical study in Section 5 shows that the BJM heuristic outperforms the BM-C heuristic. As the only difference between these heuristics is how β_{N+1} is estimated, this illustrates that the iterative approach (which requires a bit more computational work) indeed provides improved estimates of β_{N+1}.

Step 2: The distribution of the lead time demand at the central warehouse, $D_0(L_0)$, is determined according to the BM methods by treating the artificial retailer as any other retailer.

Step 3: A near optimal reorder point for the general warehouse stock, R_0, is obtained in the same way as in the BM methods. That is, by minimizing the expected holding and induced backorder costs per time unit at the central warehouse

$$\min_{R_0} \tilde{C}_0(R_0) = \frac{(h_0+\beta)}{Q_0} \sum_{y=R_0+1}^{R_0+Q_0} E_{D_0(L_0)}\left[(y - D_0(L_0))^+\right] - \beta Q\left(R_0 + \frac{Q_0+1}{2} - \mu_0\right). \quad (16)$$

It is easy to show that $\tilde{C}_0(R_0)$ is convex in R_0 and that the optimal reorder point R_0^* that minimizes (16) can be found through a simple search using the optimality condition

$$R_0^* = \max\left\{R_0 : \tilde{C}_0(R_0) - \tilde{C}_0(R_0 - 1) \le 0\right\}. \quad (17)$$

Step 4: Determination of the lead time demand at each regular retailer i, $D_i(\bar{L}_i(R_0))$ $(i = 1, \ldots, N)$, where $\bar{L}_i(R_0)$ denotes the average lead time estimated using a METRIC type approximation in accordance with the BM methods. In [7] alternative methods incorporating estimates of the lead time variance are considered, but the performance was not necessarily improved by applying these more complicated methods.

In the separate stock heuristic BM-S, the artificial retailer is treated as a regular retailer. Its reorder point $S-1$ is thus determined in Step 5 below.

In the BJM and BM-C heuristics, the lead time demand is based on a two point distribution as described in Sections 4.1 and 4.2, and the reorder point for the artificial retailer is determined using the new combined stock heuristics in Step 6 below.

Step 5: A near optimal reorder point at each retailer, R_i for $i = 1, \ldots, N$, is determined in the same way as in the BM methods by solving the fill rate constrained single-echelon problem (18). For BM-S the reorder point for the artificial retailer R_{N+1} (and thereby the critical reservation level $S_{N+1} = R_{N+1} + 1$) is also determined by solving (18) for $i = N+1$.

$$\min_{R_i} C_i(R_i) = h_i E\left[(IL_i)^+\right] = \frac{h_i}{Q_i} \sum_{y=R_i+1}^{R_i+Q_i} E_{D_i(\bar{L}_i)}\left[\left(y - D_i(\bar{L}_i)\right)^+\right]$$
$$\text{s.t.} \quad \gamma_i(\bar{L}_i(R_0), R_i) \ge \gamma_i^* \quad (18)$$

Note that for a given R_0 the reorder points can be optimized independently for each retailer.

Step 6: For the BJM and BM-C heuristics using the combined stock policy, the critical reservation level S is determined using the combined stock heuristics for compound Poisson demand (in Section 4.1) or adjusted normal demand (in Section 4.2). Note that this sixth step does not apply for the separate stock heuristic BM-S or the original BM methods.

We end this section by summarizing the different heuristics, what sets them apart and how they use the six step procedure.

BJM is defined by the full six step procedure. It uses the iterative method for estimating the induced backorder cost β_{N+1} in Step 1 and the combined stock heuristics for controlling the direct upstream demand in Step 6.

BM-C is also defined by the full six step procedure. It uses the naïve method for estimating the induced backorder cost β_{N+1} in Step 1 and the combined stock heuristics for controlling the direct upstream demand in Step 6.

BM-S is a separate stock heuristic defined by Step 1–5. It uses the naïve method for estimating the induced backorder cost β_{N+1} in Step 1, but treats the artificial retailer as a regular retailer and use Step 4 to determine S and not the combined stock heuristics.

5. Numerical study

In this section, we present results from a numerical study designed to evaluate the performances of the proposed heuristics, BM-S, BM-C and BJM, defined in Section 4 above.

The study encompasses four test series.

Test series 1 and 2 consist of 20 problems each, based on real data from two different companies.

Test series 3 consists of a complete factorial problem set of 128 researcher generated problems with compound Poisson demand.

Test series 4 is used for evaluating the alternative adjusted normal demand approximation of the BJM heuristic and consists of 32 problems taken from Test series 3.

The performances of the heuristics are evaluated using discrete event simulation models built in the ExtendSim software from Imagine That Inc. For each test problem, simulation models are used for determining the fill rates and total inventory costs associated with the policies attained with the different heuristics. All simulations are based on independent and stationary compound Poisson demand processes. The standard deviation of all fill rates estimated from simulation are below one percentage point (pp.).

Table 1
Parameter ranges for Test series 1 and 2, the central warehouse is denoted CW.

	Test series 1		Test series 2	
	Min	Max	Min	Max
Number of retailers	2	6	4	8
Target fill rate for direct upstream demand	87.8 %	95.9 %	50.0 %	99.9 %
Target fill rate for regular retailers	80.0 %	95.9 %	0.0 %	99.9 %
Fraction of direct upstream demand	14.9 %	90.8 %	0.3 %	35.4 %
Mean demand per day	0.01	4.15	0.00091	7.75
Variance to mean ratio of the demand per day	0.85	139.19	0.99	67.44
Order quantity for the cw	2	247	6	158
Order quantities for the retailers	1	37	1	92
Lead time to cw	11	73	9	9
Transportation time from cw to retailers	5	20	7	8
Holding costs per unit and day at the cw	1	1	0.1	0.1
Holding costs per unit and day at the retailers	1	1	0.1	0.1

5.1. Numerical study based on company data – Test series 1 and 2

Test series 1 and 2 are based on real data from two different companies, which are both clients of the software provider. Each test series consists of 20 problems associated with different items sold by the two companies. Test series 1 is based on data from a drilling tools company and Test series 2 on data from an electronics equipment company. The items have been chosen by the software provider to reflect the product ranges of the two companies. Other selection criteria were that the item demands should be stationary without trends, seasonality or correlations across retailers or over time. The format of the demand data we received were different for Test series 1 and 2. In case of the former, the software company performed the data analysis (because of a non-disclosure agreement with their client), and we received the empirical distributions of the customer order-sizes, together with the mean and standard deviation for the demand per day for the direct customer demand at the central warehouse, and all regular retailers. Based on this information, we fitted compound Poisson distributions with the empirical order-size distributions to the mean and standard deviation of the demand per day. These stationary and independent demand processes were then used in the simulations and in the analytical heuristics for determining the reorder levels. For Test series 2, we received more detailed demand data in the form of individual customer orders per day for each retailer, and the direct upstream demand, over a 3 year period. Analysis of the weekly demand show no significant correlation across retailers or over time. Correlation and autocorrelation coefficients are close to zero ([1^{st} quartile, median, 3^{rd} quartile] are $[-0.0405, 0.0433, 0.117]$ for the correlation coefficients, and $[-0.0359, -0.0065, 0.1137]$ for the lag 1 autocorrelation coefficients, and similar for higher time lags). Thus, confirming that the model assumption of independent and stationary demand across retailers and time are justified for the studied items. The customer order-size data was analyzed in the distribution fitting software Stat::Fit bundled with the simulation software ExtendSim. All standardized discrete distributions were rejected by the goodness of fit tests (chi-square and Kolmogorov-Smirnov tests on a 5% significance level). Therefore, as in Test series 1, the customer order-sizes are described by empirical distributions based on relative frequencies. Using these empirical distributions, we fitted compound Poisson processes to the mean and variance of the demand per day. These processes were then used in the analytical calculations and simulations,in the same way as in Test series 1.

The parameter ranges for the two test series are summarized in Table 1 where the time unit is days and the costs are in USD. Table 1 shows that the mean customer demand across the items is relatively low in both test series, although there are large differences between items. At the same time, the customer order-sizes are typically highly variable. This is further illustrated in Figs. 1 and 2 which show the mean, and the variance to mean ratio, for the customer order-sizes of the direct customer demand at the central warehouse. We can see that many items have highly variable customer order-sizes, but also that the samples include items

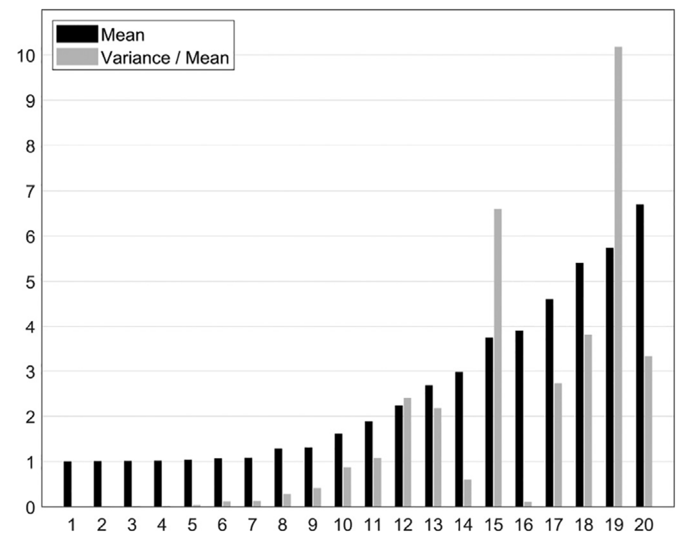

Fig. 1. Customer order-sizes for the direct customer demand at the central warehouse in Test series 1.

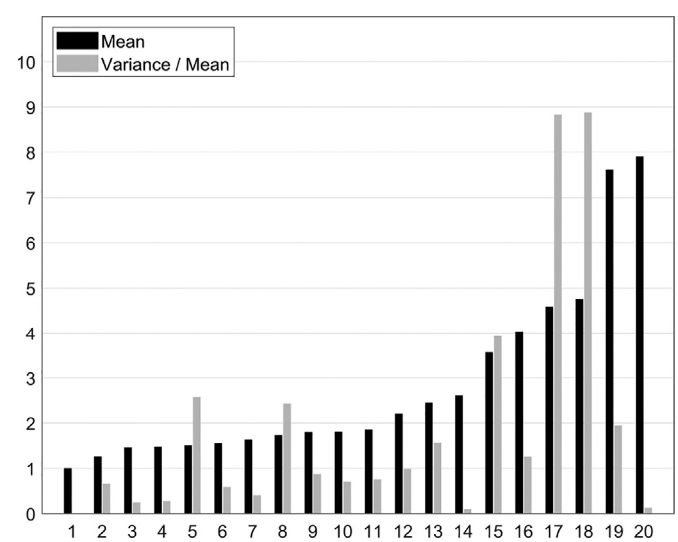

Fig. 2. Customer order-sizes for the direct customer demand at the central warehouse in Test series 2.

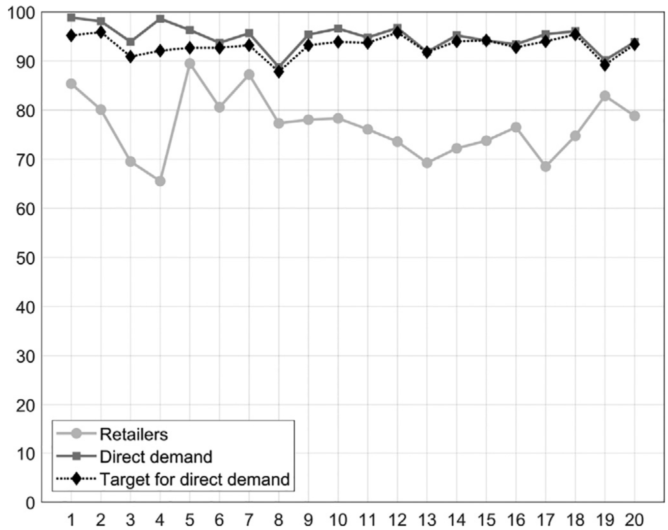

Fig. 3. Fill rates and fill rate differentiation at the central warehouse for BJM (in percent) in for the 20 problems in Test series 1.

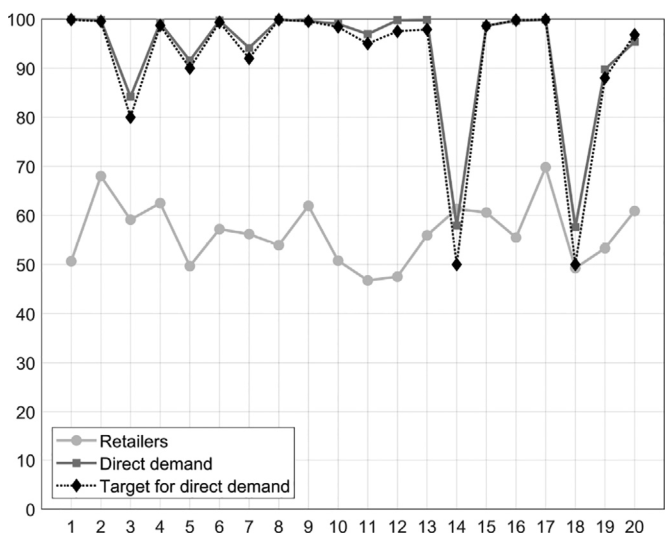

Fig. 4. Fill rates and fill rate differentiation at the central warehouse for BJM (in percent) in for the 20 problems in Test series 2.

with Poisson demand where the demand size is equal to one. The demands at the retailers show a similar structure.

When looking at Table 1 it is worth pointing out that the minimum target fill rate of 0 % for the regular retailers in Test series 2 simply indicates that an item should not be stocked at one or more locations. This occurs only in two instances where one single retailer has this target for two separate items. The second lowest target is just above 42 %. Furthermore, the minimum and maximum values of the mean demand per day along with the variance to mean ratio is taken over all customer demands, i.e., across all retailers and the direct upstream demand at the central warehouse.

For each item in the two test series, we use the BJM, BM-C and BM-S heuristics, to determine the reorder point for the general warehouse stock, R_0, the critical reservation level for the direct upstream demand, S, and the reorder points at the retailers, R_i, $i = 1, \ldots, N$. Discrete event simulation models in the ExtendSim software are then used for evaluating the fill rates and average inventory costs associated with each solution. The separate stock heuristic BM-S is used as a baseline solution.

From Figure 3 and 4 we can see that BJM renders solutions that accurately satisfies the target fill rates for the direct upstream demands for all items. Furthermore, the service at the central warehouse, is differentiated with a distinctly higher service given to the direct upstream demand than to the retailers. For two items (item 14 and item 18) in Test series 2, the fill rate target for the direct upstream demand at the central warehouse is only 50 %. At the same time, there are high fill rate targets for the demand at the regular retailers (80–92.5 % for item 14 and 50–97.7 % for item 18). This explains the small service differentiation between the two channels for these items in Figure 4. The average deviation from target fill rate across all items in Test series 1 is 1.8 pp. and in Test series 2, 1.6 pp. The corresponding values for BM-S are 3.6 pp. for Test series 1 and 5.5 pp. for Test series 2.

Table 2
Average, minimum and maximum demand weighted average deviation from target fill rates at the regular retailers for BJM (in percentage points).

	Min	Avg	Max
Test series 1	−0.36	1.67	5.32
Test series 2	−0.47	0.92	5.59

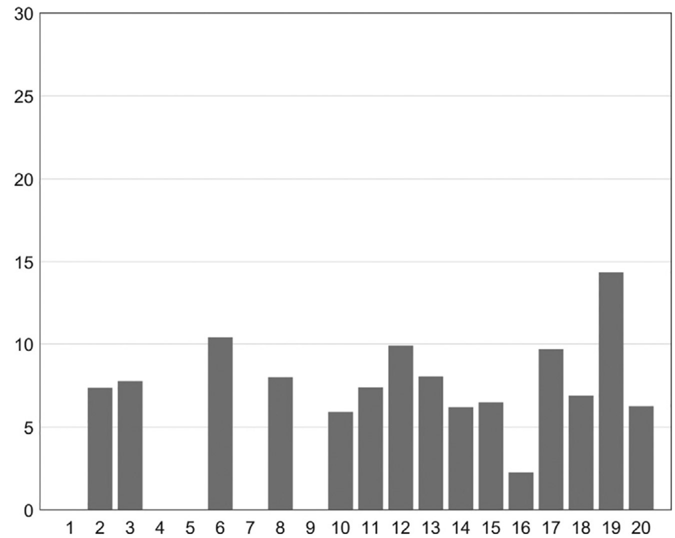

Fig. 5. Relative reduction of total inventory costs for BJM compared to BM-S (in percent) for Test series 1.

The ability for the BJM solutions to achieve the service requirements at the regular retailers is measured as the demand weighted average deviation from the given fill rate targets. Table 2 shows the average, minimum and maximum of these deviations for the two test series. We can conclude that in both Test series 1 and 2, the target fill rates for the regular retailers are satisfied on average. The corresponding values for BM-C are very similar to the ones for BJM. The ability for the original BM methods to satisfy the fill rate targets at the retailers in traditional OWMR systems has previously been evaluated in [6], showing similar results.

It is noteworthy that the BJM heuristic accurately achieves the target fill rates even though it is based on METRIC type approximations where the stochastic lead-times are approximated by their estimated means. This underestimates the lead-time variability and suggests a risk of choosing reorder levels that are too low. An explanation for why no such tendency can be seen in the numerical examples is that it is counteracted by choosing discrete reorder levels to satisfy the fill rate constraints, which typically renders fill rates slightly above target. This integer effect increases when demand is low, and a unit increase of the reorder level has a large impact on the resulting fill rate.

Figures 5 and 6 show that BJM offers significant savings in expected total inventory costs compared to the separate stock

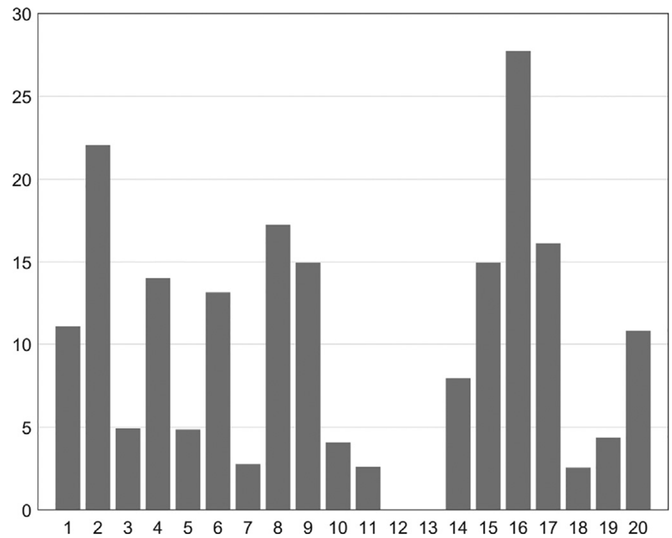

Fig. 6. Relative reduction of total inventory costs for BJM compared to BM-S (in percent) for Test series 2.

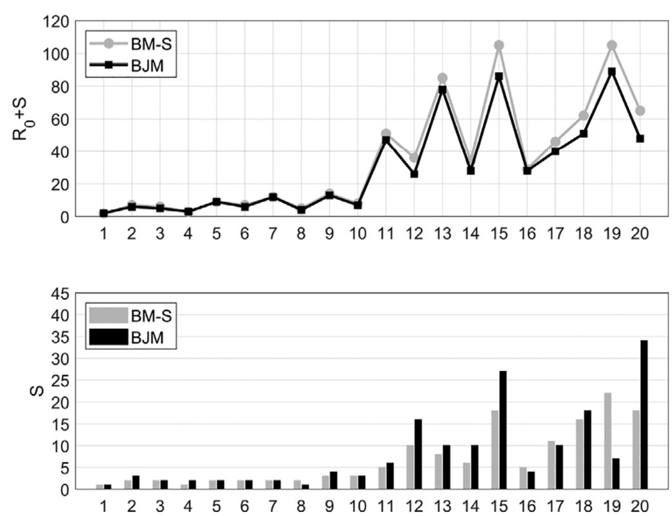

Fig. 7. Reservation level and reorder point for the total warehouse inventory for BJM and BM-S for Test series 1.

method BM-S. On average, BJM reduces the inventory costs by 5.8 % for Test series 1 and 9.7 % for Test series 2. At most, the reductions are 14.3 % and 27.7 %, respectively.

To investigate the structural differences between the combined stock solutions and the separate stock solutions, Figs. 7 and 8 show the critical reservation level S and the combined reorder point at the central warehouse $R_{cw} = R_0 + S$ obtained by BJM and BM-S, respectively.

Compared to BM-S, BJM renders a lower (or equal) reorder point for the general stock at the central warehouse, R_0. To compensate for the lower R_0, the critical reservation level, S, is generally higher for BJM to achieve the target fill rate. This leads to a larger service differentiation between the direct upstream demand and the regular retailers. At the same time, Figures 7 and 8 show that the reorder points for the combined stock at the central warehouse, $R_{cw} = R_0 + S$, for BJM is lower than or equal to BM-S, which explains why BJM requires less inventory. Although not shown in the figures, R_{cw}, for BJM is also lower than or equal to R_{cw} for BM-C.

To further investigate the BJM method's ability to determine the reservation level at the central warehouse, a simulation search has been carried out to find the value of S (for the given R_0) that renders the lowest total inventory costs. For 12 of the 20 items in Test series 1, BJM finds the best S, and for the remaining items, S is one unit too high. Consequently, all solutions provided by BJM satisfy the service constraint. For Test series 2, BJM finds the best S for 17 of the 20 items. For one item, BJM renders an S value with a fill rate 1.4 pp. below target. For the remaining two items, S is one unit too high with a fill rate 1.3 pp. and 1.9 pp. above target, respectively. Thus, for all but one of the 40 problems based on real data, the fill rate targets for the direct upstream demand are satisfied. For the problems where the provided solutions render fill rates above target, only the smallest possible adjustment of S with one unit is required to find the best value.

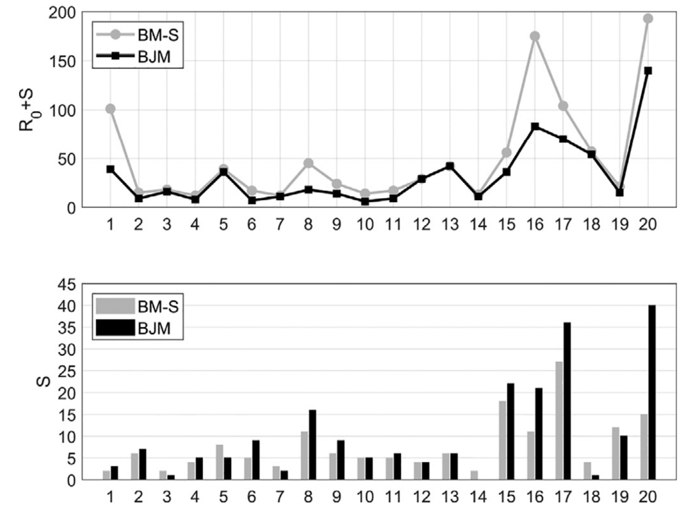

Fig. 8. Reservation level and reorder point for the total warehouse inventory for BJM and BM-S for Test series 2.

5.2. Parameter settings for Test series 3 and 4

For the researcher generated problems in Test series 3 and 4, the inventory system consists of a central warehouse with direct upstream demand and four identical retailers. The regular retailers are identical only for reasons of exposition. It is not needed from a modeling perspective, as shown by the fact that completely non-identical retailers are considered in Test series 1 and 2. The customer demand in both Test series 3 and 4 are modeled as independent compound Poisson processes, where the customer order-sizes, O_i, at retailer $i = 1, \ldots, N$ and at the central warehouse (the artificial retailer) $i = N + 1$ follows logarithmic distributions

$$f_{O_i}(d) = -\frac{\alpha_i^d}{\ln(1-\alpha_i)d} \quad i = 1, \ldots, N, N+1 \quad d = 1, 2, 3, \ldots, \quad (19)$$

where $0 < \alpha_i < 1$.

Test series 3 includes all combinations of high and low values for the seven parameters, the proportion of upstream demand, μ_{N+1}/μ_0, the coefficient of variation of the customer demand, ρ, the order quantity at the central warehouse, Q_0, and at the regular retailers, Q_i, the central warehouse lead-time, L_0, the transportation time to the regular retailers, l_i, and the target fill rate $\gamma_i^* = \gamma_{N+1}^*$, as specified in Table 3. Test series 4, used for evaluating the adjusted normal demand heuristic, consists of the 32 problems in Test series 3 with the variance to mean ratio $\rho = 5$ and the target fill rate $\gamma_i^* = 95$ %. A specification of all problem instances in Test series 3 and 4 is found in Table A.1 in Appendix A.

5.3. Results for Test series 3 – the compound Poisson demand heuristics

Table 4 summarizes the average, minimum and maximum deviation from the target fill rates of the direct upstream demand for

Table 3
Parameters used for Test series 3 and 4.

Parameter	Low value	High value
μ_{N+1}/μ_0	20 %	40 %
ρ	5	20
Q_0	20	40
Q_i	5	10
L_0	20	40
l_i	2	4
γ_i^*	95 %	99 %

BJM, BM-C and BM-S. Overall, we can see that BJM and BM-C which are based on the proposed combined stock heuristic perform much better than the separate stock heuristic BM-S, with smaller average deviations from target. Moreover, BJM has a smaller average deviation from target compared to BM-C, while still meeting the service requirements on average. The average deviation across all problems is 0.07 pp. for BJM, which is a clear improvement compared to BM-S with a total average deviation of close to 2 pp. Table 4 also shows that the maximum deviations for BJM are within one pp. from the target fill rates.

The average, minimum and maximum relative reduction in total inventory costs for BJM and BM-C compared to the baseline model BM-S are presented in Table 5. It shows that the average savings are 9.91 % for BJM and 7.18 % for BM-C. At most, the savings are 16.27 % for BJM.

Comparing BM-C with BM-S (which both use the same naïve β-etsimate) illustrates the value of using the combined stock approach instead of the separate stock approach. Looking at the values for BM-C in Table 5, we can see that the largest increase in average savings appear when increasing the variance to mean ratio of the demand per time unit, ρ, the target fill rate γ_i^*, the fraction of direct upstream demand, μ_{N+1}/μ_0 and when decreasing the order quantity at the retailers, Q_i. The explanation is that these values lead to more general stock being kept at the central warehouse which increases the importance of explicitly considering the availability of this general stock when determining S. With a higher ρ-value also follows a more variable demand, which further increases the value of considering the combined stock. This results in the largest increase of the average savings, from 5.83 % to 8.53 % when ρ increases from 5 to 20.

When comparing BJM with BM-C, we can observe larger savings of using the iterative method for determining β_{N+1} for high target fill rates and when the direct up-stream demand is a larger fraction of the total demand. The benefit of using a more correct estimate of β_{N+1} naturally increase with the weight that is put on this value when determining β in *Step 1*. The larger savings observed for higher service targets is linked to an increased overestimation of the induced backorder cost by the naïve approach. This is due to a large increase in p_{N+1}, see (15), whereas β_i is fairly insensitive to γ_i^*, see [2]. Larger savings of using BJM can also be observed for small values of ρ and to a slightly lesser extent for high values of Q_i. These shifts in parameter values result in lower β_i values. This means that overestimating β_{N+1}, which the naïve method tends to do, has a larger impact on β. In addition, a lower σ_{N+1} resulting from a lower ρ suggests a lower β_{N+1} (see, [2]). This suggests that p_{N+1} more distinctly overestimates the true β_{N+1}, which is verified by the numerical study.

Table 5
Average, minimum and maximum relative reduction in total inventory costs for BJM and BM-C compared to BM-S (in percent).

		BJM			BM-C		
		Min	Avg	Max	Min	Avg	Max
AR demand	20 %	2.57	9.03	13.68	3.93	6.83	9.43
	40 %	5.33	10.79	16.27	3.72	7.53	15.87
ρ	5	2.57	10.37	15.34	3.72	5.83	8.54
	20	6.15	9.44	16.27	6.57	8.53	15.87
Q_0	20	2.57	9.93	16.27	3.78	7.33	15.87
	40	5.38	9.89	15.34	3.72	7.04	9.43
Q_i	5	4.53	10.13	15.34	5.18	7.54	9.43
	10	2.57	9.68	16.27	3.72	6.82	15.87
L_0	20	2.57	9.43	16.27	3.77	7.32	15.87
	40	5.19	10.39	15.34	3.72	7.04	9.43
l_i	2	2.57	9.61	15.34	3.72	7.11	9.43
	4	2.57	10.20	16.27	3.77	7.25	15.87
γ_i^*	95 %	2.57	8.47	16.27	3.72	6.52	15.87
	99 %	7.51	11.35	15.34	5.72	7.85	13.63
Total		2.57	9.91	16.27	3.72	7.18	15.87

Figure 9 shows the combined reorder point, $R_{cw} = R_0 + S$, at the central warehouse obtained by the BM-S, BM-C and BJM heuristics broken down into its components, R_0 and S. Comparing BM-C with BM-S in Fig. 9, we can see that the combined stock method reduced in BM-C, reduces the reservation level S at the central warehouse, as expected. The BJM heuristic (with the iterative method for determining β_{N+1}) typically reduces the value of R_0, and to compensate for the lower general stock, the value of S needs to be increased in order to satisfy the fill rate target of the direct upstream demand. This suggests a larger service differentiation at the warehouse. At the same time, we can see from Fig. 9 that the combined reorder point, $R_{cw} = R_0 + S$, is lower for BJM than for BM-C and BM-S. This explains why the BJM solutions require less inventory.

The ability of the combined stock heuristic to find the lowest S that satisfies the fill rate target has been evaluated by simulation

Table 4
Average, minimum and maximum deviation from the fill rate targets for the direct upstream demand (in percentage points).

		BJM			BM-C			BM-S		
		Min	Avg	Max	Min	Avg	Max	Min	Avg	Max
μ_{N+1}/μ_0	20 %	−0.54	0.17	0.98	−0.38	0.27	1.16	0.73	2.00	3.79
	40 %	−0.77	−0.03	0.88	−0.07	0.36	1.33	−0.05	1,91	3.50
ρ	5	−0.77	0.08	0.98	-0.08	0.49	1.33	0.80	1.90	3.79
	20	−0.54	0.05	0.85	−0.38	0.14	0.89	−0.05	2.01	3.71
Q_0	20	−0.77	0.10	0.97	−0.22	0.35	1.33	-0.05	1.94	3.79
	40	−0.77	0.04	0.98	−0.38	0.29	1.19	0.73	1.97	3.63
Q_i	5	−0.77	0.08	0.98	−0.38	0.27	1.21	0.74	2.07	3.79
	10	−0.77	0.06	0.96	−0.33	0.36	1.33	−0.05	1.84	3.36
L_0	20	−0.50	0.14	0.98	−0.38	0.37	1.33	−0.05	1.98	3.79
	40	−0.77	−0.01	0.85	−0.33	0.27	0.93	0.73	1.93	3.69
l_i	2	−0.77	0.09	0.96	−0.33	0.32	1.31	0.75	1.98	3.69
	4	−0.77	0.05	0.98	−0.38	0.32	1.33	−0.05	1.93	3.79
γ_i^*	95 %	−0.77	0.20	0.98	−0.29	0.51	1.33	0.43	3.07	3.79
	99 %	−0.54	−0.07	0.27	−0.38	0.13	0.47	−0.05	0.84	0.96
Total		−0.77	0.07	0.98	−0.38	0.32	1.33	−0.05	1.96	3.79

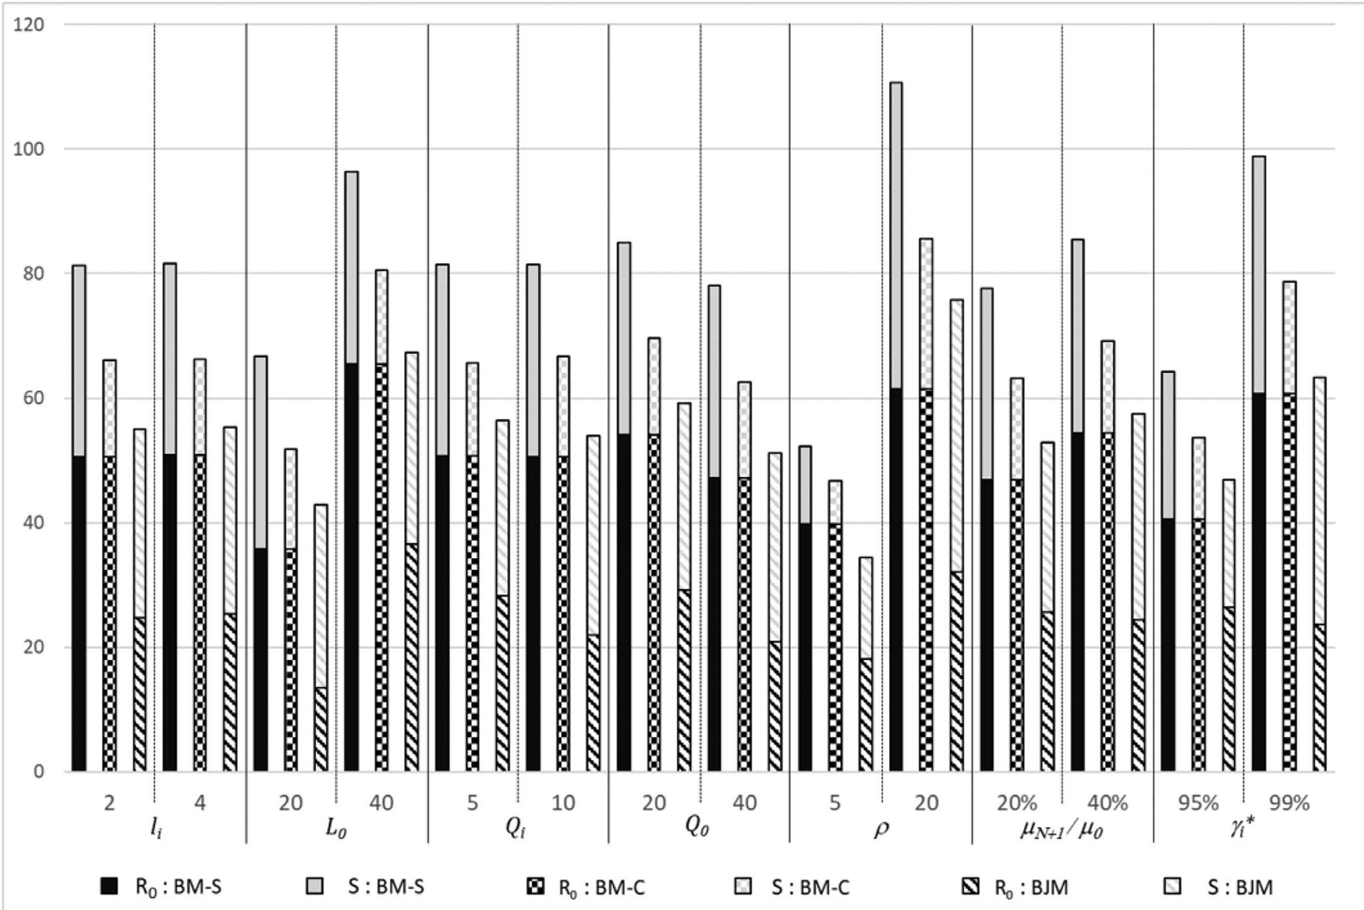

Fig. 9. Average reorder point for the general stock, R_0, reservation level, S, and $R_{cw} = R_0 + S$ at the central warehouse.

search. In 103 of the 128 test problems, BJM finds the best value of S (for the given R_0). For 11 problems, S can be reduced by one unit and still achieve the fill rate target. In 9 of the problems, the values of S need to be increased by one unit to fulfill the target service. In the remaining 5 problem instances an increase of S with 2 units is required. Consequently, in 89 % of the instances, the solution provided by BJM achieves the target fill rate (overshooting S in 7 % of the instances), and the value of S never needs to be adjusted by more than two units.

The fill rate for the general warehouse stock, serving both direct customer demand and the retailers' replenishments, is on average 60 % for the BJM heuristic. This can be compared to BM-S (and BM-C), where this fill rate is 88 %. Consequently, the BJM heuristic is achieving a larger service differentiation at the central warehouse between the direct upstream demand and the replenishment orders from the retailers.

To further investigate the quality of the BJM solutions with respect to R_0, S and R_i ($i = 1, \ldots, N$), 16 of the problems in Test series 3 have been optimized using simulation. For these problems, all parameters except γ_i^*, Q_i and l_i are alternating between the low and high values. A complete list of the chosen examples is found in Table A.1 in Appendix A. Simulation search has been used to find the optimal R_0, S and R_i values. Note that for a given R_0, BJM determines R_i, $i = 1, \ldots, N$, in the same way as in the BM methods, and the quality of these solutions are known to be good from previous work.

For one problem (no 43), the BJM method renders a fill rate for the direct upstream demand 0.68 pp below target. The reservation level S has in this case been increased by one unit to obtain a feasible solution when comparing it to the total inventory costs of the optimal solution. The relative difference in inventory costs between BJM and the optimal solutions for the 16 examples are found in Figure 10. The deviation ranges from 0 % to 6.98 % with an average value of 2.00 %. The largest deviation occurs for test problem

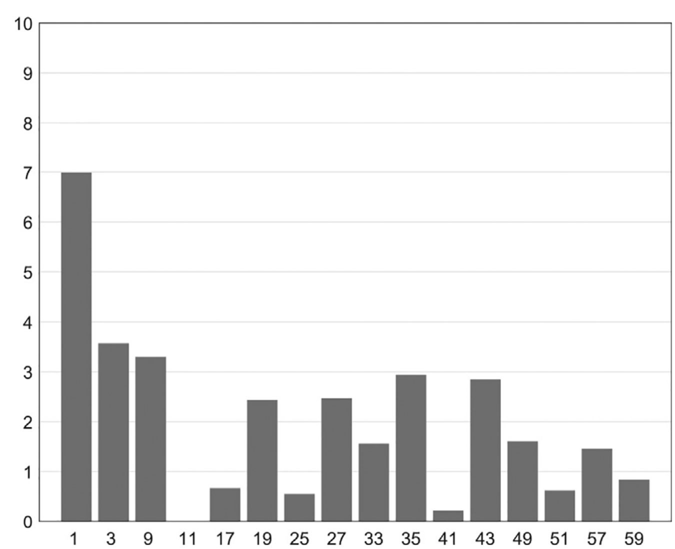

Fig. 10. *Relative difference in total inventory costs between BJM and the optimal solutions found by simulation search (in percent) for 16 problems from Test series 3. The problem instances are specified in Table A.1 in Appendix A.*

number 1, where R_0 for the optimal solution is one unit higher than the BJM solution. The larger R_0 leads to one unit lower S and R_i ($i = 1, \ldots, N$) in the optimal solution compared to the BJM solution. An adjustment of the reorder point at the retailers has a large impact on the total inventory costs, since they are identical and a reduction of R_i leads to a reduction of the inventory at all four locations. The second largest difference in total inventory costs for BJM compared to the optimal is 3.56 %.

For BJM, the deviation from the target fill rate for the direct upstream demand is on average 0.36 pp. For the optimal solutions, this deviation is on average 0.25 pp. For the retailer fill rates, the corresponding results are 0.25 pp. for BJM and 0.00 pp. for the optimal solutions. The lowest achieved fill rate for the BJM solution is 0.27 pp. below the target for the direct upstream demand and 0.15 pp. below target for the retailers. Thus, this study indicates that the BJM solutions are of good quality and close to the optimal values of R_0, S and R_i, $i = 1, \ldots, N$.

To conclude, BJM outperforms BM-S as well as BM-C with solutions that lead to fill rates closer to targets and lower total inventory costs for the system. This is achieved through a larger service differentiation at the central warehouse.

5.4. Results for Test series 4 – The adjusted normal demand heuristic

The objective with Test series 4 is to investigate the performance of the BJM heuristic for adjusted normal demand when it is used as an approximation of the compound Poisson demand model. Test series 4 encompass the 32 examples in Test series 3 with the target fill rate 95 % and lower variance to mean ratio, $\rho = 5$. Note that this means that we are testing the heuristic on quite challenging problems where the probability of negative demand in the normal distribution is rather high, and a normal demand approximation may be questioned. Nevertheless, the results show that the adjusted normal demand heuristic renders good solutions. In four examples it provides the same base-stock level, S as the compound Poisson heuristic. In one example, S is two units below, and in the remaining examples, it is just one unit below the S attained with the compound Poisson heuristic.

Figure 11 shows the deviation from the target fill rate for the direct upstream demand together with the relative difference in total inventory costs compared to the BJM heuristic with compound Poisson demand. The fill rate deviation is on average −0.68 pp. with a maximum of −1.67 pp., and the average relative difference in total inventory costs is −1.32 %. Since the adjusted normal demand heuristic tends to slightly underestimate the fill rate, the difference in total inventory costs is negative.

Finally, we note that for the problems in Test series 3 where $\rho = 20$ and $\gamma^*_{N+1} = 95$ % (not part of Test series 4), the probability of negative demand in the normal distribution is very high, and as expected the performance of the adjusted normal demand approximation therefore deteriorates. For these problems, the reservation level S is on average 6.7 units lower than for the BJM heuristic with compound Poisson demand, with an average fill rate deviation of −2.44 pp., and the relative difference in total inventory costs is −2.95 %.

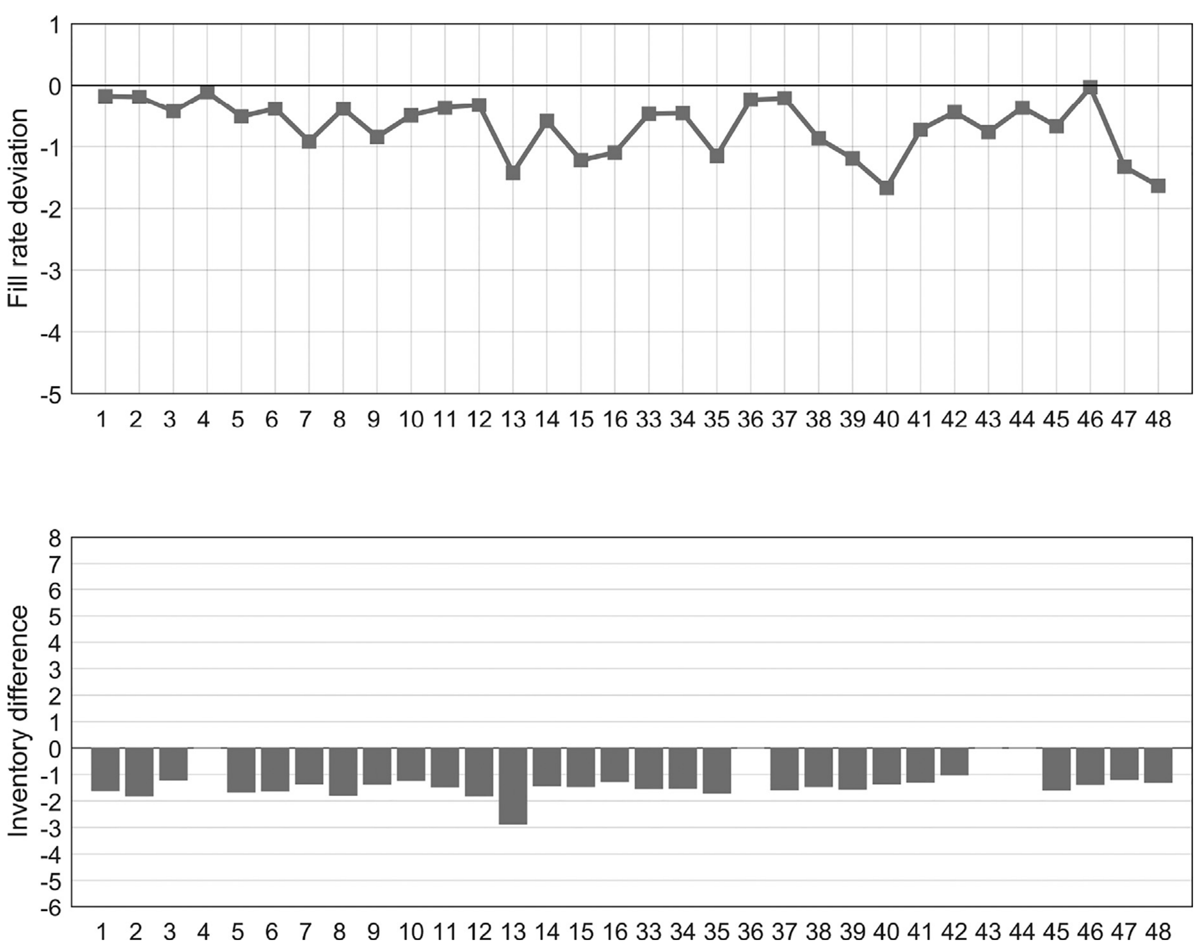

Fig. 11. Deviation from target fill rate for the direct upstream demand at the warehouse (in percentage points) and relative difference in total inventory costs for the adjusted normal demand case compared to the compound Poisson demand case (in percent) for the 48 problems in Test series 4.

6. Summary and conclusions

In this paper we have presented a combined stock approach for inventory control in real life OWMR distribution systems with direct customer demand at the central warehouse. Our work is motivated by collaboration with an inventory management software provider, and the fast growth of omni/multi-channel retailing. The real systems operated by the software provider's clients are characterized by highly variable customer order-sizes, continuous review (R, Q) policies, fill rate constraints, and direct customer demand at the central warehouse. The combined stock policy we propose represents a critical level policy at the central warehouse, which enables service differentiation between retailer replenishment orders and the direct customer demand at the central warehouse. Our technical contributions include derivations of two combined stock heuristics for determining the critical reservation level at the central warehouse; one for compound Poisson demand and one for adjusted Normal demand. We also show how these heuristics can be integrated with modified versions of the methods for control of traditional OWMR systems in [6,7]. The resulting BJM heuristics are computationally and conceptually simple enough to be used in practice, which has been an important goal for our research.

The numerical study, including both real data from two of the software provider's clients and researcher generated examples, illustrates that the proposed heuristics render near optimal solutions close to target fill rates. They also offer significant opportunities to reduce total inventory costs compared to the existing separate stock alternative. From the researcher generated examples, we can conclude that the inventory costs for the BJM solutions are on average only 2 % higher than for the optimal solutions found by simulation search. Moreover, the average deviation from target fill rates for the direct upstream demand is merely 0.07 pp. for the compound Poisson heuristic, and −0.68 pp. for the simpler adjusted normal approximation. Compared to the separate stock benchmark, BJM reduces the inventory costs by 9.9 % on average and with at most 16.3 %. For the two test series with real data, the inventory costs are reduced by 5.8 % and 9.7 % on average, and at most by 27.7 %.

CRediT authorship contribution statement

Peter Berling: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Writing – original draft, Writing – review & editing, Visualization. **Lina Johansson:** Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Writing – original draft, Writing – review & editing, Visualization. **Johan Marklund:** Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Writing – original draft, Writing – review & editing, Visualization.

Appendix A. Test series 3 and 4

All problem instances consider an inventory system with one central warehouse, four identical retailers and direct upstream demand at the central warehouse. The sum of all customer demand per time unit is $\mu_0 = \sum_{i=1}^{N+1} \mu_i = 1$. The holding costs at all locations are $h_0 = h_{N+1} = h_i = 1$ for $i = 1, \ldots, N$. The target fill rates at all locations are $\gamma_i^* = \gamma_{N+1}^* = 95$ % for $i = 1, \ldots, N$ for examples 1 to 64 and 99 % for examples 65 to 128. All other parameters are set according to Table A.1.

Test series 3 includes all 128 problem instances in Table A.1, whereas Test series 4 only includes the 32 problem instances with $\gamma_i^* = 95$ % and $\rho = 5$. The example numbers are the same in both test series. Problem instances marked with bold font in Table A.1 are optimized by simulation search, see Section 5.3.

Table A.1
Problem instances for Test series 3 and 4. Bold font marks the problem instances that have been optimized by simulation search for the compound Poisson demand case.

Examples	$\frac{\mu_{N+1}}{\mu_0}$	ρ	Q_0	Q_i	L_0	l_i	Examples	$\frac{\mu_{N+1}}{\mu_0}$	ρ	Q_0	Q_i	L_0	l_i
1, 65	20 %	5	20	5	20	2	**33**, 97	40 %	5	20	5	20	2
2, 66	20 %	5	20	5	20	4	34, 98	40 %	5	20	5	20	4
3, 67	20 %	5	20	5	40	2	**35**, 99	40 %	5	20	5	40	2
4, 68	20 %	5	20	5	40	4	36, 100	40 %	5	20	5	40	4
5, 69	20 %	5	20	10	20	2	37, 101	40 %	5	20	10	20	2
6, 70	20 %	5	20	10	20	4	38, 102	40 %	5	20	10	20	4
7, 71	20 %	5	20	10	40	2	39, 103	40 %	5	20	10	40	2
8, 72	20 %	5	20	10	40	4	40, 104	40 %	5	20	10	40	4
9, 73	20 %	5	40	5	20	2	**41**, 105	40 %	5	40	5	20	2
10, 74	20 %	5	40	5	20	4	42, 106	40 %	5	40	5	20	4
11, 75	20 %	5	40	5	40	2	**43**, 107	40 %	5	40	5	40	2
12, 76	20 %	5	40	5	40	4	44, 108	40 %	5	40	5	40	4
13, 77	20 %	5	40	10	20	2	45, 109	40 %	5	40	10	20	2
14, 78	20 %	5	40	10	20	4	46, 110	40 %	5	40	10	20	4
15, 79	20 %	5	40	10	40	2	47, 111	40 %	5	40	10	40	2
16, 80	20 %	5	40	10	40	4	48, 112	40 %	5	40	10	40	4
17, 81	20 %	20	20	5	20	2	**49**, 113	40 %	20	20	5	20	2
18, 82	20 %	20	20	5	20	4	50, 114	40 %	20	20	5	20	4
19, 83	20 %	20	20	5	40	2	**51**, 115	40 %	20	20	5	40	2
20, 84	20 %	20	20	5	40	4	52, 116	40 %	20	20	5	40	4
21, 85	20 %	20	20	10	20	2	53, 117	40 %	20	20	10	20	2
22, 86	20 %	20	20	10	20	4	54, 118	40 %	20	20	10	20	4
23, 87	20 %	20	20	10	40	2	55, 119	40 %	20	20	10	40	2
24, 88	20 %	20	20	10	40	4	56, 120	40 %	20	20	10	40	4
25, 89	20 %	20	40	5	20	2	**57**, 121	40 %	20	40	5	20	2
26, 90	20 %	20	40	5	20	4	58, 122	40 %	20	40	5	20	4
27, 91	20 %	20	40	5	40	2	**59**, 123	40 %	20	40	5	40	2
28, 92	20 %	20	40	5	40	4	60, 124	40 %	20	40	5	40	4
29, 93	20 %	20	40	10	20	2	61, 125	40 %	20	40	10	20	2
30, 94	20 %	20	40	10	20	4	62, 126	40 %	20	40	10	20	4
31, 95	20 %	20	40	10	40	2	63, 127	40 %	20	40	10	40	2
32, 96	20 %	20	40	10	40	4	64, 128	40 %	20	40	10	40	4

Appendix B. Iterative procedure to find β_{N+1}

Andersson et al. [48] show that under the assumption of normally distributed lead time demand and a constant lead time equal to the expected value $\bar{L}$, the induced backorder cost for retailer i can be computed as

$$\beta_i(\bar{L}) = (h_i + p_i)\frac{\sigma_i^2}{\mu_i Q_i}\left(\Phi\left(\frac{R_i^* + Q_i - \mu_i\bar{L}}{\sigma_i\sqrt{\bar{L}}}\right) - \Phi\left(\frac{R_i^* - \mu_i\bar{L}}{\sigma_i\sqrt{\bar{L}}}\right)\right), \tag{B.1}$$

where p_i is the backorder cost per unit and time unit at retailer i and R_i^* is the reorder point that minimizes the expected retailer costs given the assumptions above. They also show that an iterative procedure with guaranteed convergence to equilibrium solutions can be applied to determine the optimal set of induced backorder costs β_i^* $(i = 1, \dots, N)$ that coordinates the system. The method is time consuming to use for systems with many retailers. However, we use it only for finding an improved estimate of β_{N+1}, which is a fast application described by the following procedure:

Step 1 Set $\beta_{N+1} = p_{N+1}$ and determine β_i, $i = 1, \dots, N$, according to the closed form estimates in the BM methods.

Step 2 Calculate the induced backorder cost β as the weighted average of all β_i, $i = 1, \dots, N+1$, using (14).

Step 3 Determine the reorder point R_0 that minimizes (16) for the current β value.

Step 4 Estimate $\bar{L}$ with the METRIC type approximation used in the BM methods and determine R_{N+1}^* from (B.2). The latter is a sufficient optimality condition for R_{N+1} to minimize the holding and backorder costs for a continuous review single-echelon (R, Q) system with lead time $\bar{L}$ and normal demand (see [48] for details). Note, $R_{N+1} = S - 1$, $Q_{N+1} = 1$, and $G(v) = \varphi(v) - v(1 - \Phi(v))$ is the loss function.

$$\sigma_{N+1}\sqrt{\bar{L}}\left(G\left(\frac{R_{N+1} - \mu_{N+1}\bar{L}}{\sigma_{N+1}\sqrt{\bar{L}}}\right) - G\left(\frac{R_{N+1} - \mu_{N+1}\bar{L} + 1}{\sigma_{N+1}\sqrt{\bar{L}}}\right)\right) = \frac{h_{N+1}}{h_{N+1} + p_{N+1}} \tag{B.2}$$

Step 5 Update the value for β_{N+1} using (B.1). If β_{N+1} remains unchanged, or if it becomes larger than p_{N+1}, the procedure stops. Otherwise proceed to Step 6 for another iteration. Convergence is guaranteed as long as $p_{N+1} \geq h_{N+1}$ [see 48].

References

[1] Axsäter S. Supply chain operations: serial and distribution inventory systems. Handb Oper Res ManageSci 2003;11:525–59.
[2] Berling P, Marklund J. Heuristic coordination of decentralized inventory systems using induced backorder costs. Prod Oper Manage 2006;15:294–310.
[3] de Kok T, Grob C, Laumanns M, Minner S, Rambau J, Schade K. A typology and literature review on stochastic multi-echelon inventory models. Eur J Oper Res 2018;269(3):955–83.
[4] Axsäter S, Olsson F, Tydesjö P. Heuristics for handling direct upstream demand in two-echelon distribution inventory systems. Int J Prod Econ 2007;108(1–2):266–70.
[5] Axsäter S. Exact analysis of continuous review (R, Q) policies in two-echelon inventory systems with compound poisson demand. Oper Res 2000;48:686–96.
[6] Berling P, Marklund J. A model for heuristic coordination of real life distribution inventory systems with lumpy demand. Eur J Oper Res 2013;230:515–26.
[7] Berling P, Marklund J. Multi-echelon inventory control: an adjusted normal demand model for implementation in practice. Int J Prod Res 2014;52:3331–47.
[8] Gao F, Agrawal VV, Cui S. The effect of multichannel and omnichannel retailing on physical stores. Manage Sci 2021.
[9] Hu M, Xu X, Xue W, Yang Y. Demand pooling in omnichannel operations. Manage Sci 2022;68(2):883–94.
[10] Chen J, Liang Y, Shen H, Shen Z-JM, Xue M. Offline-channel planning in smart omnichannel retailing. Manuf Serv Oper Manage 2021.
[11] Bayram A, Cesaret B. Order fulfillment policies for ship-from-store implementation in omni-channel retailing. Eur J Oper Res 2021;294(3):987–1002.
[12] Jia S, Karp J, Ravi R, Tayur S. Effective online order acceptance policies for omnichannel fulfillment. Manuf Serv Oper Manage 2021.
[13] Saghiri SS, Bernon M, Bourlakis M, Wilding R. Omni-channel logistics special issue. Int J Phys Distrib.Logist. Manage. 2018;48(4):362–4.
[14] Gao F, Su X. Online and offline information for omnichannel retailing. Manuf Serv Oper Manage 2017;19(1):84–98.
[15] Gao F, Su X. Omnichannel retail operations with buy-online-and-pick-up-in-store. Manage Sci 2017;63(8):2478–92.
[16] Gallino S, Moreno A. Integration of online and offline channels in retail: the impact of sharing reliable inventory availability information. Manage Sci 2014;60(6):1434–51.
[17] Jalilipour Alishah E., Moinzadeh K., Zhou Y.-P.. Inventory fulfillment strategies for an omni-channel retailer. Working paper, University of Washington; 2015. https://ssrn.com/abstract=2659671.
[18] Govindarajan A., Sinha A., Uichanco J.. Joint inventory and fulfillment decisions for omnichannel retail networks. Ross School of Business Working Paper No. 1341; 2018. http://ssrn.com/abstract=2924850.
[19] Paterson C, Kiesmüller G, Teunter R, Glazebrook K. Inventory models with lateral transshipments: a review. Eur J Oper Res 2011;210(2):125–36.
[20] Acimovic J, Graves SC. Making better fulfillment decisions on the fly in an online retail environment. Manuf Serv Oper Manage 2014;17(1):34–51.
[21] Dai B, Chen H, Li Y, Zhang Y, Wang X, Deng Y. Inventory replenishment planning of a distribution system with storage capacity constraints and multi-channel order fulfilment. Omega 2021;102:102356.
[22] Derhami S, Montreuil B, Bau G. Assessing product availability in omnichannel retail networks in the presence of on-demand inventory transshipment and product substitution. Omega 2021;102:102315.
[23] Liu K, Zhou Y, Zhang Z. Capacitated location model with online demand pooling in a multi-channel supply chain. Eur J Oper Res 2010;207(1):218–31.
[24] Schneider F, Klabjan D. Inventory control in multi-channel retail. Eur J Oper Res 2013;227(1):101–11.
[25] Yao D-Q, Yue X, Mukhopadhyay SK, Wang Z. Strategic inventory deployment for retail and e-tail stores. Omega 2009;37(3):646–58.
[26] Agatz NAH, Fleischmann M, van Nunen JAEE. E-fulfillment and multi-channel distribution – a review. Eur J Oper Res 2008;187(2):339–56.
[27] Chiang W-yK, Monahan GE. Managing inventories in a two-echelon dual-channel supply chain. Eur J Oper Res 2005;162(2):325–41.
[28] Takahashi K, Aoi T, Hirotani D, Morikawa K. Inventory control in a two-echelon dual-channel supply chain with setup of production and delivery. Int J Prod Econ 2011;133(1):403–15.
[29] Zipkin PH. Foundations of inventory management. McGraw-Hill; 2000.
[30] Graves SC, Willems SP. Supply chain design: safety stock placement and supply chain configuration. Handb Oper Res ManageSci 2003;11:95–132.
[31] Veinott AF Jr. Optimal policy in a dynamic, single product, nonstationary inventory model with several demand classes. Oper Res 1965;13(5):761–78.
[32] Evans RV. Sales and restocking policies in a single item inventory system. Manage Sci 1968;14(7):463–72.
[33] Kaplan A. Stock rationing. Manage Sci 1969;15(5):260–7.
[34] Topkis DM. Optimal ordering and rationing policies in a nonstationary dynamic inventory model with n demand classes. Manage Sci 1968;15(3):160–76.
[35] Ha AY. Inventory rationing in a make-to-stock production system with several demand classes and lost sales. Manage Sci 1997;43(8):1093–103.
[36] Ha AY. Optimal dynamic scheduling policy for a make-to-stock production system. Oper Res 1997;45(1):42–53.
[37] Véricourt Fd, Karaesmen F, Dallery Y. Assessing the benefits of different stock-allocation policies for a make-to-stock production system. Manuf Serv Oper Manage 2001;3(2):105–21.
[38] Deshpande V, Cohen MA, Donohue K. A threshold inventory rationing policy for service-differentiated demand classes. Manage Sci 2003;49(6):683–703.
[39] Teunter RH, Haneveld WKK. Dynamic inventory rationing strategies for inventory systems with two demand classes, poisson demand and backordering. Eur J Oper Res 2008;190(1):156–78.
[40] Arslan H, Graves SC, Roemer TA. A single-product inventory model for multiple demand classes. Manage Sci 2007;53(9):1486–500.
[41] Axsäter S. Inventory control, vol 225. Springer; 2015.
[42] Sherbrooke CC. Metric: a multi-echelon technique for recoverable item control. Oper Res 1968;16(1):122–41.
[43] Sherbrooke CC. Optimal inventory modeling of systems: multi-echelon techniques, vol 72. Springer Science & Business Media; 2006.
[44] Chen F, Zheng Y-S. One-warehouse multiretailer systems with centralized stock information. Oper Res 1997;45(2):275–87.
[45] Zheng Y-S. On properties of stochastic inventory systems. Manage Sci 1992;38(1):87–103.
[46] Axsäter S. Using the deterministic EOQ formula in stochastic inventory control. Manage Sci 1996;42(6):830–4.
[47] Huh WT, Janakiraman G, Muckstadt JA, Rusmevichientong P. Asymptotic optimality of order-up-to policies in lost sales inventory systems. Manage Sci 2009;55(3):404–20.
[48] Andersson J, Axsäter S, Marklund J. Decentralized multi-echelon inventory control. Prod Oper Manage 1998;7:370–86.